

REMARKS

The Examiner is thanked for the performance of a thorough search and for the indication of allowable subject matter, namely Claims 14-19, 53, and 54 and Claims 4, 7-10, 23, 25, 26, 30, 33-36, 49, 51, and 52.

By this amendment, Claims 11-13, 27, 37-40, and 46 are amended. No claims are added or cancelled. Hence, Claims 1-56 are pending in this application.

The amendments to the claims and the new claims do not add any new matter to this application. Furthermore, the amendments to the claims were made to improve the readability and clarity of the claims and not for any reason related to patentability.

All issues raised in the Office Action mailed January 24, 2007, are addressed hereinafter.

I. SUMMARY OF THE REJECTIONS

Claims 27-52 and 55-56 stand rejected under 35 U.S.C. § 101 as allegedly being directed to non-statutory subject matter because Claims 27-52 and 55-56 recite a “machine-readable medium carrying one or more sequences of instructions” and the Specification discloses that instructions may be carried over transmission media such as light waves. Accordingly, the independent machine-readable medium claims have been amended to recite that the “machine-readable medium is one of a volatile medium or a non-volatile medium.” Therefore, reconsideration and withdrawal of the rejection of Claims 27-52 and 55-56 under 35 U.S.C. § 101 is respectfully requested.

Claims 1-3, 5-6, 27-29, 31-32 stand rejected under 35 U.S.C. § 102(e) as allegedly being anticipated by U.S. Patent Publication No. 2004/0024979 to Kaminsky et al. (“*Kaminsky*”). This rejection is respectfully traversed.

Claims 11-13 and 37-39 stand rejected under 35 U.S.C. § 103(a) as allegedly being unpatentable over U.S. Patent Publication No. 2003/0014523 to Teloh et al. (“*Teloh*”) in view of Microsoft Computer Dictionary, 5th Ed. (“*MCD*”). This rejection is respectfully traversed.

Claims 20-22, 24, 46-48, and 50 stand rejected under 35 U.S.C. § 103(a) as allegedly being unpatentable over U.S. Patent No. 6,035,379 to Raju et al. (“*Raju*”) in view of *Kaminsky*. This rejection is respectfully traversed.

II. REJECTIONS BASED ON THE CITED ART

Claims 1-3, 5-6, 27-29, 31-32 stand rejected under 35 U.S.C. § 102(e) as allegedly being anticipated by *Kaminsky*.

A. CLAIMS 1 AND 27

Claims 1 and 27 recite:

maintaining, at a first site of the plurality of sites, a record that identifies which transactions that have been executed at the first site have had their redo information replicated to the other sites of the plurality of sites; **determining a priority value associated with a transaction** that is to be performed at the first site, wherein the transaction specifies a modification to a data block; if the priority value is a first value in a set of possible values, then **committing the transaction** only after the record indicates that redo information associated with the transaction has been replicated to the other sites of the plurality of sites; and **if the priority value is a second value in said set of possible values, then committing the transaction even though the record does not indicate that redo information associated with the transaction has been replicated to the other sites of the plurality of sites.** (emphasis added)

The Office Action cites *Kaminsky* for teaching various features of Claims 1 and 27. However, there are a few fundamental differences between *Kaminsky* and Claims 1 and 27. According to Claims 1 and 27, a transaction is associated with a value that dictates when the transaction can be committed. If a transaction is associated with a first value, then the transaction is only committed after redo information associated with the transaction has been replicated to other

sites. If the transaction is associated with a second value, then the transaction is committed even though the associated redo information has not been replicated to other sites. It is well known in the art that a transaction must be committed before any modified data blocks are durably stored (i.e., on disk).

In contrast, *Kaminsky* is not concerned with when transactions commit. In fact, *Kaminsky* does not even use the terms “transaction” or “commit.” *Kaminsky* teaches a system where it is determined when file updates may be mirrored to other sites on a per-file basis. Some files may be mirrored synchronously, others asynchronously, and others not mirrored at all. In *Kaminsky*, modified data blocks are already committed (and may be **durably stored**) *before* it is determined whether a file update is to be mirrored. Thus, **there is no determination in Kaminsky of when to commit a transaction.**

To see why, paragraph [0034] and FIG. 2 of *Kaminsky* teaches that an application 60a issues a write request (i.e., an alleged transaction). File system 62a handles the request by passing data blocks representing the file update to a data I/O layer 64a, which is then sent to storage 56a. At the point where the file system 62a passes the data blocks to data I/O layer 64a, the transaction (i.e., write request) has already committed. The decision as to whether the data blocks are to be mirrored occurs at mirroring system 68a of storage 56a.

Furthermore, *Kaminsky*, at paragraph [0053], teaches: “During the time period in which the updated data blocks are being stored in the source storage system, the mirror services unit 10 consults the mirror services table 114 to determine the mirror status of the updated data blocks” (emphasis added). The Office Action cites paragraph [0054] of *Kaminsky* for teaching the last three steps of Claim 1. Paragraph [0054] discloses what actions are performed by the mirror

services unit 110, all of which occurs *after* the file is updated and, thus, *after* the corresponding transaction has committed.

Because *Kaminsky* fails to teach or suggest that a transaction may be committed at different times, *Kaminsky* fails to teach or suggest all the features of Claims 1 and 27. Therefore, Claims 1 and 27 are patentable over *Kaminsky*. Removal of the 35 U.S.C. 102(e) rejection with respect to Claims 1 and 27 is therefore respectfully requested.

B. CLAIMS 2-10 AND 28-36

Claims 2-10 and 28-36 all depend from either Claim 1 or Claim 27 discussed above and, thus, include all of the limitations of Claim 1 or Claim 27. It is therefore respectfully submitted that Claims 2-10 and 28-36 are patentable over *Kaminsky* for at least the reasons set forth herein with respect to Claims 1 and 27. Furthermore, it is respectfully submitted that Claims 2-10 and 28-36 recite additional limitations that independently render them patentable over *Kaminsky*.

C. CLAIMS 11 AND 37

Claims 11 and 37 recite:

at a first site in a plurality of sites, processing a transaction;
generating in volatile memory redo information for the processed transaction;
delaying storing the redo information to durable storage as long as (1) a data block associated with the processed transaction is not durably stored and (2) the data block is not transferred to another site of the plurality of sites; and
storing the redo information to the durable storage in response to detecting that (1) the data block is about to be durably stored or (2) the data block is about to be transferred to another site of the plurality of sites. (emphasis added)

Support for the above amendment is found at paragraph [0074] of the Specification. Thus, according to Claims 11 and 37, durably storing redo information for a transaction is delayed for as long as possible, i.e., until (1) a data block associated with the transaction is about to be

durably stored or (2) the data block is about to be transferred to another site. The cited art fails to teach or suggest **delaying**, much less delaying the durable storage of redo information.

The Office Action cites FIG. 3 and the accompanying text of *Teloh* in rejecting Claim 11. (MCD is only used to disclose volatile memory.) That portion of *Teloh* states:

The asynchronous operation of the illustrative data replication facility is as follows. An application operating on the host 16 of the local site 12 wishing to write data to the local storage device 24 first issues a write request to the local storage device 24 (step 60). The write first goes to the local data replication facility 20 operating on the host 16 of the local site 12 (step 62). The local data replication facility 20 for the local site 12 upon receipt of the write request sets a bit in a bitmap that corresponds to the data for the issued write request (step 62). The data is then written to the local storage device 24 of the local site 12 (step 64). At this point, the local data replication facility 20 copies the data into a queue to await forwarding to the remote storage device 26 of the remote site 14 (step 66).

This portion of *Teloh* describes an asynchronous approach to replicating modified data blocks to a remote site. However, *Teloh* fails to teach or suggest that redo information for a (processed) transaction is delayed as long as (1) a data block associated with the transaction is not durably stored and (2) the data block is not transferred to another site. According to paragraph [0074] of Applicant's Specification, “[t]he steps illustrated in FIG. 7 may be used to advantageously delay replication of a write operation to a redo log. It is advantageous for a database server to reduce the latency of writing data to a redo log.”

Based on the foregoing, the combination of *Teloh* and MCD fail to teach or suggest all the features of Claims 11 and 37. Accordingly, Claims 11 and 37 are patentable over *Teloh* and MCD. Removal of the 35 U.S.C. § 103(a) rejection with respect to Claims 11 and 37 is therefore respectfully requested.

D. CLAIMS 12-13 AND 38-39

Claims 12-13 and 38-39 all depend from either Claim 11 or Claim 37 discussed above and, thus, include all of the limitations of Claim 11 or Claim 37. It is therefore respectfully submitted that Claims 12-13 and 38-39 are patentable over *Teloh* for at least the reasons set forth herein with respect to Claims 11 and 37. Furthermore, it is respectfully submitted that Claims 12-13 and 38-39 recite additional limitations that independently render them patentable over *Teloh*.

E. CLAIMS 20 AND 46

Claims 20 and 46 recite a method or machine-readable medium for mirroring data that comprises:

at the first site, durably storing a data block **prior to** durably storing redo information about changes made to the data block; and
at the first site, durably storing the redo information after the changes have been replicated to the other sites in the plurality of sites. (emphasis added)

According to Claim 20, two events must occur before redo information corresponding to a data block is durably stored at a first site. One of those events is durably storing the data block at the first site, and the other event is the changes made to the data block are replicated to other sites. As will be shown hereafter, *Raju* fails to teach or suggest that the above two events occur before redo information is durably stored at the first site.

In rejecting Claim 20, the Office Action cites two portions of *Raju*: (1) a portion of the Background section and (2) a portion that describes a conventional approach to logging changes. However, these portions of *Raju* teach away from the features of Claim 20 by employing conventional techniques for storing redo information.

The first portion of *Raju* above merely states:

In transaction processing, a "transaction" is a logical unit of work that is to be atomically performed. For example, a transfer of funds between bank accounts constitutes a single transaction that entails the two operations of debiting from one account and crediting the other account. Transaction processing guarantees that if a transaction executes some updates and then a failure occurs before normal termination is reached, the updates are undone. A transaction either executes in its entirety or is totally canceled.

This first portion is cited to show that "a transaction can include updates to multiple databases" (Office Action, page 9). According to this portion, a transaction may include multiple operations. Also, this portion merely teaches that "a 'transaction' is a logical unit of work that is to be atomically performed." If a transaction includes multiple write operations and one of the write operations is unsuccessful, then any other write operations that have executed are canceled and the transaction is not committed.

The second portion of *Raju* states:

FIG. 5 is a flowchart that shows the steps that are performed to log update records into the log file 30. The log file 30 is located in a persistent secondary storage, such as a disk storage. Initially, a transaction is logged by writing the update records for the operations of the transaction to the log file 30 (step 42 in FIG. 5). The associated operations of the transaction are then performed (step 44 in FIG. 5). When the transaction is committed, a log record indicating commitment of the transaction is added to the log file 30 (step 46 in FIG. 5).

This second portion of *Raju* is cited to show that "a transaction is performed (step 44) prior to storing a record in a commit log (step 46)" (Office Action, page 9). Performing a transaction does not mean that any modifications the transaction made are durably stored. If FIG. 5 of *Raju* illustrated when the data blocks (that are affected by the operations in step 44) are durably stored, then that step would follow step 46. This is made clear in col. 3, lines 35-38, where it states: "NTFS does this because it does not know whether the modifications were flushed to a disk in time before the failure, despite the transactions being committed" (emphasis added).

Indeed, “FIG. 5 is a flowchart illustrating the steps that are performed by a conventional system that supports NTFS” (col. 4, lines 45–46) and “NTFS supports a write-ahead logging approach” (col. 2, lines 59–60). In a system using write-ahead logging (WAL), all modifications are written to a log before the modifications are written to a database. In other words, according to WAL (and *Raju*), redo information is durably stored *before* a corresponding data block is durably stored. However, Claim 20 teaches the opposite. According to Claim 20, the redo information is durably stored *after* the data block is durably stored.

Furthermore, the Office Action cites the same portions of *Raju* for allegedly disclosing “at the first site, durably storing the redo information *after* the changes have been replicated” (emphasis added), as Claim 20 recites. However, *Raju* also fails to teach or suggest this ordering. *Raju* is silent as to replication, much less replicating changes that have been made to a data block.

The Office Action cites *Kaminsky* for disclosing a plurality of sites (page 9). The Office Action then asserts that “*Raju* discloses that his transaction system is able to send the same update to multiple databases, such as bank accounts” (page 9). It is respectfully submitted that this statement is incorrect. The first cited portion of *Raju* above states that two update operations are performed (i.e., a debit operation and a credit operation), not one update. Furthermore, *Raju* merely states that “a transfer of funds between bank accounts constitutes a single transaction that entails the two operations of debiting from one account and crediting the other account.” It cannot be extrapolated from the above statement that the same update is sent to multiple databases. However, regardless of whether the same update is sent to multiple databases, the combination of *Raju* and *Kaminsky* still fails to teach or suggest that changes to a data block are

replicated to other databases or sites **before** the corresponding redo information is durably stored at the first site.

Because *Raju* teaches away from Claims 20 and 46, and *Kaminsky* is only used to disclose a plurality of sites, Claims 20 and 46 are patentable over the art of record. Therefore, removal of the 35 U.S.C. § 103(a) rejection with respect to Claims 20 and 46 is respectfully requested.

F. CLAIMS 21-24 AND 47-50

Claims 21-24 and 47-50 all depend from either Claim 20 or Claim 46 discussed above and, thus, include all of the limitations of Claim 20 or Claim 46. It is therefore respectfully submitted that Claims 21-24 and 47-50 are patentable over *Raju* and *Kaminsky* for at least the reasons set forth herein with respect to Claims 20 and 46. Furthermore, it is respectfully submitted that Claims 21-24 and 47-50 recite additional limitations that independently render them patentable over *Raju* and *Kaminsky*.

III. CONCLUSION

It is respectfully submitted that all of the pending claims are in condition for allowance and the issuance of a notice of allowance is respectfully requested. If there are any additional charges, please charge them to Deposit Account No. 50-1302.

The Examiner is invited to contact the undersigned by telephone if the Examiner believes that such contact would be helpful in furthering the prosecution of this application.

Respectfully submitted,

HICKMAN PALERMO TRUONG & BECKER LLP

/DanielDLedesma/

Daniel D. Ledesma

Reg. No. 57,181

Date: April 24, 2007

2055 Gateway Place, Suite 550
San Jose, CA 95110-1089
Telephone: (408) 414-1229
Facsimile: (408) 414-1076